



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/665,062	09/17/2003	Roger J. Bamford	50277-2277	5488
29989	7590	10/05/2006	EXAMINER	
HICKMAN PALERMO TRUONG & BECKER, LLP 2055 GATEWAY PLACE SUITE 550 SAN JOSE, CA 95110			MIZRAHI, DIANE D	
			ART UNIT	PAPER NUMBER
			2165	

DATE MAILED: 10/05/2006

Please find below and/or attached an Office communication concerning this application or proceeding.



UNITED STATES DEPARTMENT OF COMMERCE
U.S. Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450

APPLICATION NO./ CONTROL NO.	FILING DATE	FIRST NAMED INVENTOR / PATENT IN REEXAMINATION	ATTORNEY DOCKET NO.
---------------------------------	-------------	---	---------------------

10/665,062

EXAMINER

ART UNIT PAPER

20060930

DATE MAILED:

Please find below and/or attached an Office communication concerning this application or proceeding.

Commissioner for Patents

Please see attached correspondence: Examiner's Amendment which formally corrects Applicant's record and Interview Summary

Diane Mizrahi
PRIMARY EXAMINER

EXAMINER'S AMENDMENT

An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with Attorney Brian Hickman 9-27-2006.

The application has been amended as follows:

Claims:

1. (Previously presented) A method for managing data, comprising:
maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;
assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;
when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item while said particular data item continues to reside at said particular location;

reassigning ownership of the particular data item from the particular node to another node without moving the particular data item from said particular location on said persistent storage;

after the reassignment, when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to said other node for the other node to perform the operation on the particular data item while said particular data item continues to reside at said particular location,

wherein:

said persistent storage is a first persistent storage of a plurality of persistent storages used by a multi-node database system,
the method further comprises reassigning ownership of a second data item from a first node that has access to said first persistent storage to a second node that has access to a second persistent storage but does not have access to said first persistent storage, and

the method further comprises reassigning ownership of the second data item by moving the second data item from said first persistent storage to said second persistent storage.

2. (Original) The method of Claim 1 wherein the step of reassigning ownership of the particular data item from the particular node to another node includes updating an ownership map that is shared among the plurality of nodes.

3. (Currently Amended) ~~The method of Claim 5, wherein the plurality of nodes are nodes of a multi-node database system~~ A method for managing data, comprising:

maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;

assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;

when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;

while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node;

after the reassignment, when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to said other node for the other node to perform the operation on the particular data item,

wherein:

said persistent storage is a first persistent storage of a plurality of persistent storages used by said multi-node database system; and

the method further comprises reassigning ownership of a second data item from a first node that has access to said first persistent storage to a second node that has access to a second persistent storage but does not have access to said first persistent storage; and
wherein the step of reassigning ownership of the second data item includes moving the second data item from said first persistent storage to said second persistent storage.

4. (Currently Amended) The method of Claim 3, wherein the plurality of nodes are nodes of a multi-node database system ~~the multi-node database system includes nodes that do not have access to said first persistent storage.~~

5. (Currently Amended) The method of Claim 4 wherein the multi-node database system includes nodes that do not have access to said first persistent storage A method for managing data, comprising:

~~maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;~~
~~assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;~~
~~when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the~~

~~particular node for the particular node to perform the operation on the particular data item;~~

~~while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node;~~

~~after the reassignment, when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to said other node for the other node to perform the operation on the particular data item,~~

~~wherein:~~

~~said persistent storage is a first persistent storage of a plurality of persistent storages used by said multi-node database system; and~~
~~the method further comprises reassigning ownership of a second data item from a first node that has access to said first persistent storage to a second node that has access to a second persistent storage but does not have access to said first persistent storage; and~~

~~wherein the step of reassigning ownership of the second data item includes moving the second data item from said first persistent storage to said second persistent storage.~~

6. (Currently Amended) The method of Claim 3 5 wherein the step of reassigning ownership of the particular data item from the particular node to another node is

performed in response to the addition of said other node to said multi-node database system.

7. (Currently Amended) The method of Claim 3 5 wherein:
the step of reassigning ownership of the particular data item from the particular node to another node is performed in anticipation of the removal of said particular node from said multi-node database system; and
the method further comprises the step of, in anticipation of the removal of said particular node from said multi-node database system, physically moving from said persistent storage to another persistent storage a second data item that is reassigned from said particular node to a node of said multi-node database system that does not have access to said persistent storage.8. (Original) The method of Claim 3 5 wherein the step of reassigning ownership of the particular data item from the particular node to another node is performed as part of a gradual transfer of ownership from said particular node to one or more other nodes.
9. (Original) The method of Claim 8 wherein the gradual transfer is initiated in response to detecting that said particular node is overworked relative to one or more other nodes in said multi-node database system.
10. (Previously Presented) A method for managing data, comprising:

maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage,
wherein the plurality of nodes are nodes of a multi-node database system;
assigning exclusive ownership of each of the persistent data items to one of the nodes,
wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;
when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;
while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node;
wherein the step of reassigning ownership of the particular data item from the particular node to said other node is performed as part of a gradual transfer of ownership from said particular node to one or more other nodes in said multi-node database system,
wherein the gradual transfer is initiated in response to detecting that said particular node is overworked relative to the one or more other nodes,
wherein the gradual transfer is terminated in response to detecting that said particular node is now longer overworked relative to the one or more other nodes ; and

after the reassignment, when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to said other node for the other node to perform the operation on the particular data item.

11. (Previously Presented) A method for managing data, comprising:
 - maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage,
 - wherein the plurality of nodes are nodes of a multi-node database system;
 - assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;
 - when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;
 - while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node;
 - wherein the step of reassigning ownership of the particular data item from the particular node to another node is performed as part of a gradual transfer of ownership to said other node by one or more other nodes, wherein said gradual transfer is

initiated in response to detecting that said other node is underworked relative to the one or more other nodes in said multi-node database system; after the reassignment, when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to said other node for the other node to perform the operation on the particular data item.

12. (Previously Presented) A method for managing data, comprising :
maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;
assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;
when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data; while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node; after the reassignment, when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the

operation to said other node for the other node to perform the operation on the particular data item; and
after a first node has been removed from the multi-node system, continuing to have a set of data items owned by the first node.

13. (Previously Presented) The method of Claim 12, further comprising :
reassigning ownership of a data item from the first node to a second node in response to detecting that the workload of said second node has fallen below a predetermined threshold.

14. (Original) The method of Claim 1 wherein:
at the time said particular data item is to be reassigned to said other node, the particular node stores a dirty version of said particular data item in volatile memory; and the step of reassigning ownership of the particular data item from the particular node to another node includes writing said dirty version of said particular data item to said persistent storage.

15. (Previously Presented) A method for managing data, comprising:
maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;
assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;

when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;

while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node,

wherein:

at the time said particular data item is to be reassigned to said other node, the particular node stores a dirty version of said particular data item in volatile memory; and

the step of reassigning ownership of the particular data item from the particular node to another node includes forcing to persistent storage one or more redo records associated with said dirty version, and purging said dirty version from said volatile memory without writing said dirty version of said particular data item to said persistent storage; and

said other node reconstructs said dirty version by applying said one or more redo records to the version of the particular data item that resides on said persistent storage.

16. (Previously Presented) A method for managing data, comprising:
maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;

assigning exclusive ownership of each of the persistent data items to one of the nodes,

wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;

when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;

while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node,

wherein:

at the time said particular data item is to be reassigned to said other node, the particular node stores a dirty version of said particular data item in volatile memory; and

the method further includes the step of transferring the dirty version of said particular data item from volatile memory associated with said particular node to volatile memory associated with said other node.

17. (Original) The method of Claim 16 wherein the step of transferring the dirty version is performed proactively by the particular node without the other node requesting the dirty version.

18. (Original) The method of Claim 16 wherein the step of transferring the dirty version is performed by the particular node in response to a request for the dirty version from said other node.

19. (Previously Presented) A method for managing data, comprising:

maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;

assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;

when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;

reassigning ownership of the particular data item from the particular node to another node,

wherein:

the step of reassigning ownership of the particular data item from the particular node to another node is performed without waiting for a transaction that is modifying the data item to commit;

the transaction makes a first set of modifications while the particular data item is owned by the particular node; and

the transaction makes a second set of modifications while the particular data item is owned by said other node.

Art Unit: 2165

20. (Original) The method of Claim 19 further comprising rolling back changes made by said transaction by rolling back the second set of modifications based on undo records in an undo log associated with said other node, and rolling back the first set of modifications based on undo records in an undo log associated with said particular node.
21. (Previously Presented) A method for managing data, comprising :
maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes, the persistent data items including a particular data item stored at a particular location on said persistent storage;
assigning exclusive ownership of each of the persistent data items to one of the nodes, wherein a particular node of said plurality of nodes is assigned exclusive ownership of said particular data item;
when any node wants an operation performed that involves said particular data item, the node that desires the operation to be performed ships the operation to the particular node for the particular node to perform the operation on the particular data item;
while the particular node continues to operate, reassigning ownership of the particular data item from the particular node to another node;
the other node receiving a request to update said data item;
determining whether the particular node held exclusive-mode or shared-mode access to the data item; and
if the particular node did not hold exclusive-mode or shared-mode access to the data item, then the other node updating the particular data item without waiting for the

particular node to flush any dirty version of the data item, or redo for the dirty version, to persistent storage.

22. (Previously Presented) The method of Claim 1 further comprising the steps of:
in response to transferring ownership of said particular data item to said other node,
aborting an in-progress operation that involves said particular data item;
after ownership of the particular data item has been transferred to said other node, re-executing the in-progress operation.
23. (Original) The method of Claim 1 wherein:
an operation that involves said particular data item is in-progress at the time the transfer of ownership of said particular data item is to be performed;
the method further includes the step of determining whether to wait for said in-progress operation to complete based on a set of one or more factors; and
if it is determined to not wait for said in-progress operation to complete, aborting said in-progress operation.
24. (Original) The method of Claim 23 wherein said set of one or more factors includes how much work has already been performed by said in-progress operation.
25. (Previously Presented) A method of managing data, the method comprising the steps of:
maintaining a plurality of persistent data items on persistent storage accessible to a plurality of nodes;
assigning ownership of each of the persistent data items to one of the nodes by

assigning each data item to one of a plurality of buckets by enumerating individual data-item-to-bucket relationships; and
assigning each bucket to one of the plurality of nodes by enumerating individual bucket-to-node relationships;

wherein the node to which a bucket is assigned is established to be owner of all data items assigned to the bucket;

when a first node wants an operation performed that involves a data item owned by a second node, the first node ships the operation to the second node for the second node to perform the operation.

26. (Original) The method of Claim 25 wherein the step of assigning each data item to one of a plurality of buckets is performed by applying a hash function to a name associated with each data item.
27. (Original) The method of Claim 25 wherein the step of assigning each bucket to one of the plurality of nodes is performed by applying a hash function to an identifier associated with each bucket.
28. (Original) The method of Claim 25 wherein the step of assigning each data item to one of a plurality of buckets is performed using range-based partitioning.
29. (Original) The method of Claim 25 wherein the step of assigning each bucket to one of the plurality of nodes is performed using range-based partitioning.
- 30-31. (Cancelled).

32. (Original) The method of Claim 25 wherein the number of buckets is greater than the number of nodes, and the bucket-to-node relationship is a many-to-one relationship.
33. (Original) The method of Claim 25 further comprising the step of reassigning from a first node to a second node ownership of all data items that are mapped to a bucket by modifying a bucket-to-node mapping without modifying a data-item-to-bucket mapping.
34. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 1.
35. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 2.
36. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 3.
37. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 4.
38. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 5.

39. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 6.
40. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 7.
41. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 8.
42. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 9.
43. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 10.

44. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 11.
45. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 12.
46. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 13.
47. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 14.
48. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 15.
49. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 16.

50. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 17.
51. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 18.
52. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 19.
53. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 20.
54. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 21.
55. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 22.

56. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 23.
57. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 24.
58. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 25.
59. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 26.
60. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 27.
61. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 28.

62. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 29.

63-64. (Cancelled).

65. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 32.

66. (Original) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 33.

67. (Original) A method for use in a multi-node shared-nothing database system, the method comprising the steps of:
a first node of said multi-node shared-nothing database system initially functioning as exclusive owner of a first data item and a second data item, wherein said first data item and said second data item are persistently stored data items within a database managed by the multi-node shared-nothing database system;
without changing the location of a first data item on persistent storage or shutting down said first node, reassigning ownership of the first data item from the first node to a second node of said multi-node shared-nothing database system; and

after reassigning ownership, the first node continuing to operate as the owner of the second data item, and to handle all requests for operations on said second data item.

68. (Previously Presented) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 67.
69. (Previously Presented) The method of Claim 12, further comprising: reassigning ownership of data items from the first node to one or more other nodes in response to detecting requests for operations that involve said data items.
70. (Previously Presented) A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 69.

Comments

The date on which the petition under 37 CFR 1.136(a) and the appropriate extension fee have been filed is the date for purposes of determining the period of extension and the corresponding amount of the fee. In no case may an applicant reply outside the SIX (6) MONTH statutory period or obtain an extension for more than FIVE (5) MONTHS beyond the date for reply set forth in an Office action. A fully responsive reply must be timely filed to avoid abandonment of this application.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

As allowable subject matter has been indicated, Applicant's response must either comply with all formal requirements or specifically traverse each requirement not complied with. See 37 CRF 1.111(b) and MPEP section 707.07(a).

Other Prior Art Made of Record

The prior art made of record and not relied upon is considered pertinent to Applicant's disclosure. U.S. patents and U.S. patent application publications will not be supplied with Office actions. Examiners advises the Applicant that the cited U.S. patents and patent application publications are available for download via the Office's PAIR. As an alternate source, all U.S. patents and patent application publications are available on the USPTO web site (www.uspto.gov), from the Office of Public Records and from commercial sources. For the use of the Office's PAIR system, Applicants may refer to the Electronic Business Center (EBC) at <http://www.uspto.gov/ebc/index.html> or 1-866-217-9197.

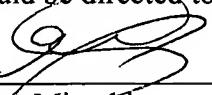
Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Diane D. Mizrahi whose telephone number is 571-272-4079. The examiner can normally be reached on Monday-Thursday.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeffrey Gaffin can be reached on (571) 272-4146. The fax phone numbers for the

organization where this application or proceeding is assigned are (703) 872-9306 for regular communications and (703) 305-3900 for After Final communications.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.



Diane Mizrahi
Primary Patent Examiner
Technology Center 2100

August 29, 2006